

AD-A126 563

INTEROPERABILITY ANALYSIS(U) SYSTEM DEVELOPMENT CORP
SANTA MONICA CA C M SWITZKY 08 JAN 82
SDC-TM-7038/220/00 DCA100-80-C-0044

1/1

UNCLASSIFIED

F/G 17/2


NL

END

FORMED

1980

2010



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

SDC

System Development Corporation

2500 Colorado Avenue, Santa Monica, CA 90406. Telephone (213) 820-4111

series

base no./vol./reissue

7038/220/00

author *Carl M. Switzky*
Carl M. Switzky

technical *David J. Kaufman*
David J. Kaufman

release *David J. Kaufman*
David J. Kaufman

for Charles A. Savant

date 1/8/82

TM

a working paper

This document was produced by
System Development Corporation in performance of Contract DCA100-80-
C-0044

DCEC PROTOCOLS STANDARDIZATION PROGRAM INTEROPERABILITY ANALYSIS

ABSTRACT

This report presents an analysis of various methods for interconnecting the Integrated AUTODIN System (IAS) to public data networks. The approaches to interoperability include (1) requiring an IAS phaseover to an X.25 based network, and then incorporating X.75 interfaces to other networks; (2) providing an interface between the DoD protocols at IAS sites and the X.25 based networks; and (3) providing a translation between the protocol services offered by the two networks.

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

83 04 07 036

AD A 126563

DTIC FILE COPY

DTIC
SELECTED
APR 7 1983
H

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 7038/220/00	2. GOVT ACCESSION NO. AD-A146563	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Interoperability Analysis		5. TYPE OF REPORT & PERIOD COVERED interim technical report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Carl M. Switzky		8. CONTRACT OR GRANT NUMBER(s) DCA100-80-C-0044
9. PERFORMING ORGANIZATION NAME AND ADDRESS System Development Corporation 2500 Colorado Ave. Santa Monica, CA 90406		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS P.E. 33126K Task 1053.558
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Communications Engineering Center Switched Networks Engineering Directorate 1860 Wiehle Ave., Reston, VA 22090		12. REPORT DATE 8 Jan 82
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) N/A		13. NUMBER OF PAGES 37
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE N/A
16. DISTRIBUTION STATEMENT (of this Report) Approved for Public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) N/A		
18. SUPPLEMENTARY NOTES This document represents results of interim studies which are continuing at the DCEC of DCA.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Protocols, Data Communications, Data Networks, Protocol Standardization, Interoperability		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report presents an analysis of various methods for interconnecting the Integrated AUTODIN System (IAS) to public data networks. The approaches to interoperability include (1) requiring an IAS phaseover to an X.25 based network, and then incorporating X.75 interfaces to other networks; (2) providing an interface between the DoD protocols at IAS sites and the X.25 based networks; and (3) providing a translation between the protocol services offered by the two networks.		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

CONTENTS

1. INTRODUCTION.....	1
2. PROTOCOL ARCHITECTURE LAYERING.....	2
2.1 Integrated AUTODIN System Architecture.....	2
2.2 Non-DoD Architecture.....	4
2.3 Architectural Differences and Deficiencies.....	5
3. INTEROPERABILITY ALTERNATIVES.....	7
3.1 IAS Phaseover to X.25 Based Networks.....	7
3.2 X.25 Network Backbone.....	10
3.3 Protocol Translation.....	12
4. TRADEOFF ANALYSIS.....	15
4.1 Feature Analysis.....	16
4.1.1 Phaseover Approach 16	
4.1.2 X.25 Backbone Approach 17	
4.1.3 Protocol Translation Approach 18	
4.2 Technical Complexity.....	18
4.2.1 Phaseover Approach 18	
4.2.2 X.25 Backbone Approach 20	
4.2.3 Protocol Translation Approach 21	
4.3 Economic Analysis.....	29
4.3.1 Phaseover Approach 29	
4.3.2 X.25 Backbone Approach 30	
4.3.3 Protocol Translation Approach 30	
4.4 Performance Analysis.....	30
4.4.1 Phaseover Approach 30	
4.4.2 X.25 Backbone Approach 31	
4.4.3 Protocol Translation Approach 32	
4.5 Survivability Characteristics.....	32
4.5.1 Phaseover Approach 33	
4.5.2 X.25 Backbone Approach 33	
4.5.3 Protocol Translation Approach 33	
4.6 Security Impact.....	33
4.6.1 Phaseover Approach 33	
4.6.2 X.25 Backbone Approach 34	
4.6.3 Protocol Translation Approach 34	
4.7 Generality for Growth.....	34
4.7.1 Phaseover Approach 34	
4.7.2 X.25 Backbone Approach 35	
4.7.3 Protocol Translation Approach 35	
4.8 Transition Impact.....	35
4.8.1 Phaseover Approach 35	
4.8.2 X.25 Backbone Approach 36	
4.8.3 Protocol Translation Approach 36	
5. CONCLUSIONS.....	36
6. REFERENCES.....	37

Accession For	
NTIS GRA&I	
DTIC TAB	
Unannounced	
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	



LIST OF FIGURES

<u>Figure</u>		<u>Page</u>
1.	IAS Protocol Architecture.....	3
2.	Non-DoD Protocol Architecture.....	6
3.	X.25 Based IAS Approach.....	8
4.	X.25 Network Backbone Approach.....	11
5.	X.25 to DoD Protocol Translation Approach.....	13
6.	Protocol Translator Layered Architecture.....	14

1. INTRODUCTION

The Integrated AUTODIN System (IAS) is being developed to provide distributed computing capabilities for the Department of Defense in the 1980s. The IAS will consist of host computers, user terminals, workstations, and other specialized processing equipment that are connected together via AUTODIN II and other network backbones[2]. To facilitate effective communication between the IAS components, standardized DoD protocols will be implemented in accordance with defined layering principles.

Parallel to the development of DoD protocol standards has been the development of recommendations for network interface standards by the International Telegraph and Telephone Consultative Committee (CCITT)[1]. These recommendations are part of an overall protocol architecture that is intended to fit into the ISO Reference Model[4]. These standards are intended to allow development of network systems for use by domestic non-DoD government agencies, government agencies in other countries, and by domestic and foreign commercial communications vendors. Public Data Networks (PDNs) such as TELENET or TYMNET are examples of network systems that have been developed to provide network service to these non-DoD users. Other examples of PDNs following CCITT recommendations may be found within NATO countries.

PDNs may be compared to IAS network backbones with regard to the type and extent of service offered. They both are intended to act as the communications media between connected hosts and users by implementing protocols up to the network layer of the ISO reference model. Protocols for transport, session, presentation, and application layers must be implemented in hosts, front-ends, or other processors connected to the PDN or backbone.

The IAS network backbones and PDNs implement their network communication capabilities in inherently different manners. DoD networks support a packet transfer service of individual data packets from one site to another. A datagram service of this type has no requirement for connection establishment. PDNs support the X.25 virtual call mechanism. This mechanism requires that a circuit be established between sites that wish to exchange data. A datagram service for PDNs has been approved by CCITT and incorporated into the X.25 recommendation. This facility, which allows the PDNs to provide service similar to the datagram service of the DoD backbones, has had only limited support by PDN suppliers.

Potential DoD interoperability with PDN facilities is motivated both by survivability requirements, i.e. PDNs could provide additional redundant communication capabilities for the DoD in crisis situations or when other IAS backbones have disrupted service, and by the desire to support connections to non-DoD sites. The ability to communicate with non-DoD sites would allow data exchange with NATO facilities, for example. In addition, the proposed adoption of X.25 as a federal standard, may necessitate DoD interoperability with PDNs for communication with other federal agencies within the United States.

This report evaluates three approaches for interconnecting the IAS with PDN sites. The evaluation will be based on analyzing the approaches in terms of the features obtained, the technical complexity, the economic costs, the

impact on performance, the survivability characteristics, the security impact, the generality for growth, and the impact during a transition period. The three approaches are (1) phasing the DoD architecture to a CCITT compatible one so that a uniform architecture exists; (2) utilize PDNs as carriers of DoD traffic by having sites implement an interface between the DoD Internet Datagram Protocol, IP and the X.25 virtual call service; and (3) translate between the protocols used in the two architectures.

The remainder of this report is organized into four sections. As background to the discussion of the alternative approaches, Section 2 describes the protocol layering used in the two network architectures and the differences between them. Next, the alternative approaches for interconnection are discussed to give an idea of what is involved in each. Section 4 presents the trade-offs for each of the approaches based on the criteria listed above. The final section draws conclusions from the analysis.

2. PROTOCOL ARCHITECTURE LAYERING

Interoperability refers to providing users of one network the capability of accessing services on another network. Within each network, users have access to services via a set of layered protocols. Interoperability will thus be achieved by developing an interface between the two networks at suitable protocol levels, or by implementing a common protocol architecture on both networks.

The DoD and X.25-based architectures addressed by this report contain facilities for providing connections between different sites, and connections across different networks. In addition, the architectures include facilities for terminal control from both ends of a connection. These facilities can be mapped onto the network, transport and presentation layers in the ISO Reference Model. Neither of the protocol architectures, however, is complete; they are both lacking established protocols that correspond to some of the layers. In addition, there is no commonly accepted non-DoD protocol layering above the X.25 level. For the purposes of this report, assumptions will be made about potential non-DoD protocol layerings to produce an example non-DoD architecture. This section will describe the two architectures and point out differences between them and deficiencies that impact interoperability.

2.1 Integrated AUTODIN System Architecture

The purpose of the IAS protocol architecture is to standardize the techniques for providing and obtaining IAS services. This is achieved through the use of a layered, modular architecture. Figure 1 illustrates the structure and identifies the major DoD protocols. The purpose and description of each of the layers of this architecture are discussed in [5].

At the base of the architecture are protocols necessary to provide the interface to the backbone network carrier being used. Examples of network interface protocols used by DoD are the Segment Interface Protocol (SIP) and the




Figure 1. IAS Protocol Architecture

1822 Host-IMP interface specification. These interfaces provide the host site the means of getting data in and out of the attached network.

The lowest common element of the protocol architecture is the Internet Protocol (IP) which provides a datagram-like service and internetting capabilities. This protocol provides the ability to send data packets to other sites on the network based on the address in the packet. When the address indicates that the packet is destined to another network, IP determines which site on the network is acting as a gateway to that network and sends the packet to that site.

The basic packet-transfer service of IP is used to provide more sophisticated services to the upper protocol layers. Currently the only commonly used service is the network connection offered by the Transmission Control Protocol (TCP), although other transport protocols might be utilized to meet additional requirements such as voice, where throughput and delay characteristics outweigh reliability concerns.

TCP provides reliable, sequenced, and flow controlled transfer of data between two ends of a connection. These features are achieved by coordination of the TCP modules through the exchange of connection related control information. TCP service also includes the opening and closing of connections, as well as out-of-band signaling on the connections.

The DoD Session Protocol will provide the control of presentation services that require multiple connections, quarantine service, and other interaction management functions. The exact services of this protocol are currently under development. The Telnet/THP protocols provide the means by which a user and/or application physically access the IAS services. These protocols support the concept of a network virtual terminal so that it is not necessary for application programs to be cognizant of the characteristics or the particular type of terminal being used. The protocols perform the mapping between the network virtual terminal format used below the protocols and the format necessary for proper control of the terminal. The protocols also provide for control of the terminal characteristics from each end of the connection.

2.2 Non-DoD Architecture

As previously mentioned, there is no commonly accepted protocol layering above the X.25 interface at PDN sites. The PDN suppliers provide network access to subscriber hosts and terminals; they do not specify how the network is used by the hosts above the network access level. There is no formal definition of the overall protocol architecture. Other standards groups such as ISO, ANSI, and ECMA are developing standards for protocols above the X.25 level that may be adopted by hosts. The eventual adoption of a set of these standards will form the non-DoD architecture.

The CCITT interface separates equipment into two classes. Data Circuit-Terminating Equipment (DCEs) provide access to the communication capabilities of the network. Data Terminal Equipment (DTEs) provide the computing capabilities of a site. PDNs provide and support DCEs and all underlying hardware and communications equipment; subscriber hosts to PDNs are considered to be DTEs.

With recommendation X.25, CCITT has specified the interface between a subscriber DTE and the PDN's DCE. This recommendation divides the interface into three levels: the physical, the link, and the packet.

The physical level specifies the electrical interface between the DTE and the DCE.

The link level specifies the control of the data link formed by the physical interface between the DTE and the DCE. This control includes starting and stopping the link, packet delimiting, and error control.

The packet level uses the link to provide multiple channels; X.25 specifies the different types of packets that can be sent and received on these channels. PDNs may offer permanent virtual circuits, virtual calls, and datagram services. Through the different type of packets, the DTEs can access these

services. The virtual call service consists of the procedures for establishing calls, controlling the flow of data between the DTE and the DCE, sending interrupts, and closing the calls.

Transferring packets between networks is performed at an X.75 interface. Packets intended for another network are routed within the network to the interface.

The datagram facility of X.25 allows for the transmission of data and a response without the establishment of a virtual call. At this time there seems to be little support for this service by the PDNs.

ECMA and ANSI have defined a proposed standard for a transport level protocol[3]. This protocol is a candidate for providing transport service above the X.25 level for PDN sites. This protocol provides for five different classes of service: Simple, Basic, Flow Control, Error Recovery, and Error Detection and Recovery. These services will be provided on an end-to-end basis. Another candidate for a transport protocol is one proposed by NBS which offers services that are fundamentally similar to those of TCP.

Like the DoD architecture, CCITT also supports the concept of a network virtual terminal through the X.3/X.28/X.29 interfaces. These recommendations specify the terminal characteristics (X.3), the local user interface to the terminal characteristics (X.28), and the remote process interface to the terminal characteristics (X.29). These interfaces are intended to sit directly on top of X.25, but for the purposes of analysis, will be used as a model for a terminal control protocol and therefore separated by other layers. This model is illustrated in Figure 2.

2.3 Architectural Differences and Deficiencies

At this point it is useful to identify some of the philosophical and technical differences between the two network architectures so that the problems of interconnecting between them can be seen.

The first philosophical difference is in the nature of the environment of the networks and their intended applications. DoD networks are viewed as vehicles for exchanging and distributing classified intelligence information under both routine and crisis situations. As such, the architecture must incorporate mechanisms for controlling access to data, for handling disrupted service, and for allowing interruptions for higher priority operations. These concerns are not incorporated into the PDNs because their function is limited to providing network service in what is presumably a benign environment. The PDNs therefore have different standards of service to meet. Users who require enhancements of available services must provide these enhancements in their own software.

Another philosophical difference between the two architectures is in the nature of their development. The DoD work stemmed from the original NCP and TELNET protocols which were implemented in hosts using DoD networks. Development of these protocols through research and their use by the network




Figure 2. Non-DoD Protocol Architecture

community has led to the current versions of IP, TCP, and THP. The experience gained from these protocols, which fit into the ISO network, transport, and presentation layers, has been incorporated into the DoD architecture. The objective of CCITT was to develop a standard by which different users could interface to developing commercial networks in a manner similar to telephone switching systems. This led to emphasis on the X.25 interface (network layer) and left the design of the upper level protocols and services to the users, and the design of the network internals to the PDN suppliers. CCITT also perceived the need to standardize the terminal interface, which led to the X.3/X.28/X.29 recommendations for the presentation layer.

This difference in development objectives naturally leads to a difference in the primitive mechanisms employed in the two architectures. As mentioned earlier, the DoD architecture is based on a datagram service while the non-DoD architecture is based on a circuit service. For DoD services that do not use the connection oriented TCP, interoperability will be difficult to provide in the PDN environment.

The virtual call provided by an X.25 based network is different from the TCP connection because many characteristics of the call are provided on a hop-by-hop rather than on an end-to-end basis. Each of the X.25 and X.75 interfaces and the hops within the PDN that make up a call provide the desired characteristics individually. Proposed non-DoD transport protocols will provide service that is in the same category as TCP, but the developing state of these protocols, and the uncertainty of which ones will be adopted, will complicate the interoperability problem.

The two primitive mechanisms also lead to different approaches to internetting between different networks. The DoD architecture calls for gateway sites that are able to transfer the datagrams from one network to another without necessity for established connections. PDNs use the X.75 interface between the networks so that the X.25 circuits can span multiple networks. Thus, the internetting capability in the CCITT architecture is provided by and under the control of, the network supplier; in the DoD architecture it is performed at specific sites and is not part of the service provided by a network. Another difference between the two internetting approaches is that IP allows for alternate internetwork paths without reestablishing the connection; the disruption of an X.25/X.75 virtual circuit requires reestablishment of the circuit.

An issue related to internetting is the address formats that are used in the two network types. X.25 uses a 14 decimal based address format, while the DoD uses a 32 bit based address format. Approaches to achieve interoperability must include the capability to translate between these two formats.

A final deficiency in both network types is the lack of established session level protocols. A non-DoD version of such a protocol would not address the DoD concerns for security, access control, and multiple user sessions. The different approaches taken will lead to additional difficulties in obtaining services for one network type from the other. Lack of concrete session level information requires that analysis of the impact of this layer on interoperability be an issue for further study.

3. INTEROPERABILITY ALTERNATIVES

Three basic approaches have been suggested to obtain interoperability between the IAS and PDNs. Each of these alternatives provide different capabilities with different costs. Each of three approaches is described in the following subsections.

3.1 IAS Phaseover to X.25 Based Networks

A common protocol architecture used by both PDN and DoD networks eliminates most of the problems associated with interoperability. Since the DoD is not able to impose its architecture on PDNs, this approach will consider a phaseover of the DoD architecture to one based on X.25 networks. The approach, which uses the X.75 recommendation for internetting, is illustrated in Figure 3.

The X.75 interface used in this approach provides interoperability by supporting virtual calls between the DoD and non-DoD networks. Unless sites on both networks use compatible protocols above the X.25 level, the interoperability will be limited. Differences between TCP and the ANSI transport protocol, between the yet-to-be-defined session protocols, and between TELNET/THP and X.3/X.28/X.29 mechanisms will prevent any interoperability between DoD and non-DoD networks even though X.25 circuits can be established between them.




Figure 3. X.25 Based IAS Approach

The use of common upper level protocols must also be considered as part of the phaseover approach if the X.25 network is to be more than an alternate network backbone for the DoD. Use of X.25 networks as alternate backbones will be considered in the next section (X.25 Network Backbone).

To interoperate with the upper level protocols at non-DoD sites, the DoD sites must support the compatible protocols. This support of non-DoD protocols may be done on an IAS-wide basis, or it may be done only at selected sites that need the capability. If only a limited number of sites require this capability, then it may be reasonable for only those sites to support the necessary protocols. In this case, other sites would only be able to interoperate with the non-DoD sites through a protocol translation mechanism. Use of this translation mechanism will be considered in Section 3.3 (Protocol Translation). The approach discussed in this section will be concerned with the phaseover on an IAS-wide basis to protocols compatible with the those used at non-DoD sites.

The utilization of non-DoD protocols constitutes a change in the direction that has been taken by the DoD in the past. Unless there is support by the PDNs for the datagram facility of X.25, the most significant aspect of this change is the elimination of the datagram service. The datagram service is implemented as the Internetwork Protocol (IP).

IP includes in its design support for a number of primitives that can be used to provide many different types of transport service. These primitives include options for indicating precedence, reliability, speed, resource trade-off, source routing, stream identification, time-stamping, and security

marking. With these primitives, transport level protocols can be built that emphasize reliable delivery of data, prompt delivery of data, separation of data streams, security marking of data streams, routing of data over specific paths, delivery of data to multiple sites, and delivery of data without a connection requirement. Transport services other than the reliable, sequenced, flow control type that has been developed by ANSI and ECMA may not be supportable in the X.25 environment.

The preferred direction of the phaseover approach is to develop protocols which are compatible with the protocols being used on other networks but provide the features that will meet DoD requirements. Thus TCP would be replaced by the adopted transport protocol, and other transport protocols would be developed making use of X.25 options to provide the necessary services. This approach might require an enhanced X.25 network to be developed by the DoD. This network would support all the PDN X.25 options as well as placing increased emphasis on the X.25 datagram option so that the other transport protocols could be developed. With this approach, interoperability with non-DoD sites on other PDNs would be possible with the adopted transport protocol, and the other transport protocols would provide the services required for the DoD-unique applications.

Because of various factors including the investment in existing applications, protocols, and hardware by the DoD, the immature state of X.25 based networks, and the undeveloped nature of the higher level protocols, part of the phaseover approach must include a staging plan by which elements can be smoothly integrated into the existing DoD system. This staging must take into account the utilization that can be made of existing X.25 networks, the length of time necessary to develop and test an enhanced DoD X.25 network (if this is deemed necessary), the availability of standardized protocols, and the development schedule of future DoD protocols.

The first stage of a phaseover might remove IP from the architecture by building an interface between TCP and X.25. The objective of this phase would be to use TCP as the transport protocol, X.25 as the network interface, and X.75 for internetting. This will allow interoperability between sites using TCP and the replacement of the DoD hardware with a vendor-supplied X.25 network. This stage would allow TCP and other existing protocols above to remain unchanged.

The interface would coordinate TCP connections with X.25 virtual calls, and provide the translation of addresses between DoD and CCITT formats. Each time a TCP connection was to be established to a remote site on either the same or other network, the interface would first establish an X.25 virtual call to the other site, making use of X.75 interfaces if necessary. When the TCP connection was closed, the X.25 virtual call would also be closed. This will require the interface to be aware of the occurrence of these TCP actions so they can be translated into the appropriate X.25 actions.

As an alternative strategy for this first stage, X.25 could be used as a medium between hosts without explicit coordination between TCP connections and X.25 virtual calls. An X.25 virtual call would be used as a site-to-site conduit on which TCP connections would be multiplexed. This strategy would be

more consistent with the TCP/IP interface because it would not require a virtual call to be established explicitly for each TCP connection. It would be inconsistent with the objective of phasing to the non-DoD transport/network type of interface in which this coordination exists.

After the use of PDNs as network backbones is established, the next stage of the phaseover would replace TCP with the transport protocol adopted by other PDN users. This replacement will extend the PDN interoperability to sites that support the same transport protocol and should simplify the interface between the transport and X.25 layers. A new interface will have to be designed to interface the higher level protocols to the new transport protocol.

Eventually, all the IAS protocols would be replaced by ones that are compatible with those adopted by PDN users, including the session and presentation layer protocols. Specifically, TELNET or THP would be replaced by X.3/X.28/X.29. The application programs using these services would have to be adapted to new protocols.

The early utilization of X.25 based networks would allow the development, testing, and integration of an enhanced X.25 network that would provide services necessary to meet DoD requirements for security, reliability, and priority. Once an enhanced X.25 network is developed it could be substituted for the standard X.25 network with a minimum of service interruption for existing transport protocol applications. New transport protocols and related applications could then be installed to make use of the enhanced backbone while compatibility was maintained with non-DoD networks.

3.2 X.25 Network Backbone

The complete phaseover of the IAS to X.25 based networks is an extensive move to provide interoperability. A less extensive move is to devise a method by which sites that support DoD protocols can operate on X.25 type networks and require PDN sites with interoperability requirements to support the DoD protocols. This approach would also allow for alternative communication paths between DoD sites with X.25 access by using the PDNs as additional network backbones.

In order for DoD protocols to operate on an X.25 based network, an interface must be developed that will allow IP to transfer datagrams on the X.25 network. Interoperability between sites on a PDN and sites on a DoD backbone would be achieved by an IP gateway site that would transfer the internet datagram packets between the two network types. This is illustrated in Figure 4.

This approach is less extensive than the previous approach because the IP-to-X.25 interface need only be implemented at sites that will use PDN type networks for communications. This includes DoD sites that act as PDN gateways, DoD sites that require additional backbone communication capability, and other non-DoD sites that have DoD communication requirements. This approach does require that all sites that are to use X.25 networks support the DoD protocols




Figure 4. X.25 Network Backbone Approach

from the IP level up.

The IP-to-X.25 interface must convert between the DoD and CCITT addressing formats and must use the X.25 mechanisms for sending the packets to the appropriate destination site. When the source site is on the PDN, and the packet is intended for a site on a DoD backbone, the IP would choose the gateway site on the PDN and cause the packet to be sent to that site. The gateway site would transfer the packet into the other network.

Because the X.25 virtual call is the most uniformly supported mechanism, it would be the likely candidate for transferring packets between sites. One approach for utilizing the virtual calls is to establish them as needed, and to clear them when they no longer seem to be used. This approach will require the interface to maintain a table of the currently established virtual calls. Whenever a packet was received, the table would be checked to determine if a virtual call exists to the destination site chosen for the packet. If a call exists to that site, the packet would be packaged as X.25 data and sent on the

existing circuit; otherwise, the virtual call would have to be established first. Virtual calls would be cleared whenever the activity level on them drops, indicating that the transport level connections that required the circuit had probably been closed. If connections that require the destination site were still open, then the next packet to that site would cause the interface to reestablish the virtual call.

These virtual calls would form a conduit between sites and not be explicitly associated with specific transport level connections or their opening or closing. In this way they, the interface could mimic the type of service that is provided by a datagram network. Transport level connections would be aware of little difference in the network behavior, with the exception of the additional delay whenever it is necessary to establish the X.25 connection.

A second approach is possible if the number of sites on the network is small. When this is the case it would be possible to use the X.25 Permanent Virtual Circuit facility with circuits always established between each of the sites. In addition to eliminating the need for maintaining the table necessary for establishing and clearing virtual calls, this would also eliminate the delay in sending a packet whenever a virtual call had to be established.

A third approach is possible if the destination sites for the X.25 packets is always on the same PDN and that PDN supports the datagram facility. When this is the case, the interface could use this facility to transfer the packet without the need for establishing a virtual call or using an existing permanent virtual circuit. If the X.25 virtual calls that make use of an X.75 interface between two PDNs are established, and the other network does not support the datagram facility, then this approach could not be used.

3.3 Protocol Translation

The third interoperability approach is to translate between the protocols used on the two network types. This approach would allow sites that do not support the DoD protocols to interoperate with the DoD sites and also minimize the impact on the IAS. The configuration for protocol translation is illustrated in Figure 5.

The figure illustrates an independent site that functions both as a gateway between the networks and as the location at which the translation between the network sites occurs. These two functions can also be divided between two different sites if a means of connecting the protocol translation site to the gateway site is provided. This discussion will be concerned with the translation aspects of the approach, and not the physical location of the translator. Translation is independent of the location of the translator.

The general idea of protocol translation is to provide interoperability through a composite connection formed by two separate connections and bridged by the translator. Performing the bridging between the connection halves requires the translation of the protocols used in each of the network types. This translation must convert between the services offered by the protocols on both network types. To the extent that the translation is complete, and all




Figure 5. X.25 to DoD Protocol Translation Approach

services offered by one network are available from the other, interoperability will be achieved. Shortcomings of the translation will result in the reduction of the type and quality of service that users of one network type can obtain from the other.

The use by sites on both network types of similar layered protocol architectures suggests a similar layered structure for a protocol translator. The layered structure implies that the translator would consist of pairs of protocol modules. Each of the modules of a pair would act as the logical end of its protocol for its network type. The transferring of protocol features across the bridge would be done by a translation module that interfaces the pair of protocol modules. There may be no need for the translation of some

protocols when features offered by those protocols are intended to take advantage of services that are local to the network type. These pairs of protocol modules, and necessary translation modules, would be structured in a layered architecture as is done at the host sites.

In general, a protocol provides its services by encapsulating and unwrapping data it receives from upper and lower protocol levels. When data from an upper level is encapsulated, the protocol information is added to the message. Data that is unwrapped from a lower level is done so with the interpretation of the protocol information that was added to the message by the sending protocol. Other non-protocol related information is considered data by the protocol.

Modules in the translator must act by interpreting the protocol information received from lower levels, requesting comparable services from the associated protocol module, and passing data to the next higher level protocol. This is illustrated in Figure 6.




Figure 6 -- Protocol Translator Layered Architecture

Not all features offered will be directly passed across the translation interface for a specific protocol because of the necessity of coordinating some features with the specific data. These features may be incorporated into the data that is passed to the upper level and then turned back into the protocol service when the data reaches the associated protocol module.

So far, this discussion has centered on the internals of the protocol translator. Another general issue related to achieving interoperability via protocol translation is the external view of the translator. The purpose of this translator is to allow communication between sites on different networks. Connections over a single network or a set of similar networks generally attempt to make the network transparent to the user -- they try to provide service that is equivalent to the service the user would receive if directly connected to the host. To achieve this transparency, it is necessary for the user and/or host to inform the network of the network handling procedures of certain network related information. This information includes indication of the destination and characteristics of the connection support desired. This type of information indicates an underlying awareness of the network nature of the communication that is to take place.

When two dissimilar network types are to be interconnected, the question arises as to how much of the interconnecting process the ends of the connections should be aware of. At one extreme, one connection end should only be aware its network type, as in the single network type case. The translator should be smart enough to handle all processing necessary to support the connection on the other network type. At the other extreme, just as some awareness of the underlying nature of the network is necessary for single network type usage, awareness of the two network types may be necessary for using both of them together.

The first approach leads to an ideal solution to the problem because it does not require additional processing at either end of the connection; the extra processing is done at the translator. The second approach allows the translator to be less sophisticated but requires additional protocols at both ends of the connection for control of the bridge between the two connection halves.

The first approach may be difficult to implement completely because of the necessity of some explicit control over the operation of the translator from connection ends. The second approach may also present difficulties because of the trouble of embedding new protocols to handle this interoperation in the existing protocol structure. The feasible solution will lie somewhere between these two extremes.

4. TRADEOFF ANALYSIS

The evaluation of the various alternatives for interconnection must be based on a comparison of the time and expenditure costs versus the capabilities obtained. The following criteria will be used to analyze the alternative approaches for interconnecting AUTODIN II with public data networks:

1. Feature Analysis - What each approach will and will not do.
2. Technical Complexity - How difficult it will be to utilize the approach in terms of the technical problems that need to be solved.
3. Economic Analysis - A comparison of the costs that will be encountered as the result of the need to develop and maintain custom hardware/software, the ability to use commercially available hardware/software, and the amount of such hardware/software that must be used in an operational system.
4. Performance Analysis - What impact of each of the approaches on throughput and performance can be expected.
5. Survivability Characteristics - How well a system based on each approach can be expected to handle disruption of communication capabilities.
6. Security Impact - How effectively end-to-end security can be integrated into the system proposed by each of the approaches.
7. Generality for Growth - How each of the approaches will be able to cope with the addition of new services and technologies.
8. Transition Impact - What effect the incorporation of the additional hardware/software will have on the existing IAS operation.

In the following sections the three different approaches will be discussed with respect to each of the criteria.

4.1 Feature Analysis

Each of the alternatives will provide a different type of service for interoperability. This section will compare the features and deficiencies of the approaches. These characteristics will also be evaluated with respect to the objectives of interoperability discussed in the Introduction (Section 1). These objectives included the ability to interoperate with sites that do not utilize DoD protocols, and to provide additional communication capabilities over PDNs for DoD sites during crisis situations.

4.1.1 Phaseover Approach

The phaseover approach allows the interconnection between sites on connected networks through the X.25 virtual call mechanism and X.75 interfaces. Because this approach entails converting IAS sites to an architecture that is supported by PDNs, it will potentially result in compatible protocol architectures existing at DoD and non-DoD sites. This approach will meet both interoperability objectives.

The protocol architecture used by non-DoD sites does not include the datagram service as the common base for all upper level protocols, and therefore

removes from DoD use all services that are dependent on datagram service. These include broadcast, real time, and transaction oriented services. This approach is also dependent on X.75 interfaces being implemented and working on the networks that are to be utilized for interoperability.

Other features provided by IP but not supported by X.25 virtual calls include precedence treatment for services requiring high priority, a method for guaranteed allocation of resources to insure a specified level of service, source control of the network path to be used, reporting of the path taken, and the capability to label the security compartment of traffic.

As pointed out in the discussion of this approach, it may be possible to develop an enhanced version of an X.25 based network that would also provide services necessary to meet DoD requirements. This form of the network would maintain compatibility with PDNs, and therefore be capable of interoperating with non-DoD sites. However, if an enhanced network were adopted, it would mean that the DoD would be unable to use PDNs as alternative backbones because they would not provide the appropriate level of service.

This approach constitutes a change in the basic philosophy that has been pursued in the development of the DoD architecture. As such, its implementation would have to be performed in a phased approach so that each change could be tested and adjusted to. General interoperability between DoD and non-DoD hosts would not occur until DoD sites supported protocols compatible with those adopted by the non-DoD sites. Therefore, the availability of the DoD and non-DoD interoperability in a practical manner would be dependent on both DoD phasing schedules, and the schedule by which non-DoD sites agree on and adopt their protocols.

4.1.2 X.25 Backbone Approach

The X.25 network backbone approach meets the interoperability objective of utilizing a PDN as an alternative network backbone. Because it requires that the sites being connected support the DoD protocols, it is unable to provide any interoperability with non-DoD sites.

With this approach, internetting may be performed by the DoD IP protocol or by X.75 interfaces. IP can only request the delivery of datagrams to sites it considers to be on the same network. If the X.75 interfaces works properly, then sites on other PDNs may be considered to be on the same network by an IP module.

If an IP module is restricted to sending packets to sites on its local PDN, the IP/X.25 interface can make use of special features of the local PDN. For example, a datagram service supported by the network would improve performance and simplify the interface. Otherwise, the IP module will not know if such a service is available on other connected PDNs.

Using an X.25 based network may restrict the nature of the service that an IP module can offer. Real time support may be impractical with the additional overhead of the X.25 virtual call mechanism.

4.1.3 Protocol Translation Approach

The translation approach allows for the connection between IAS and PDN sites without the need for transition to a common protocol architecture. It meets the interoperability objective of providing communication between DoD and non-DoD sites with a minimum impact on the existing architectures. Protocol translation would be implemented at specific sites that may also act as the gateways to the PDNs. Depending on the details of the implementation, impact on other DoD sites would be minimized and should be restricted to sites that actually have interoperability requirements.

The value of protocol translation is limited by the number of protocols that can be translated, the compatibility that can be achieved by the translation, and the performance that can be maintained. To make use of an application service in one network type from another network type, all lower level protocols must be translated. The translation between protocols will cause the loss of some of the protocol services because of the lack of equivalent services. The larger the number of protocols that must be translated, the larger the loss of service, and therefore the more difficult it will become to use the application, thus reducing the degree of interoperability.

Protocol translation does not meet the objective of providing an alternative network backbone because it does not provide a method of connecting to DoD hosts without performing the translation. Two sites could be connected if translations were performed in both directions, but the reduction in service and performance using this approach would likely be unacceptable.

4.2 Technical Complexity

In order for each of the approaches to support their features, problems associated with the approaches must be solved. The evaluation of the approaches with respect to the technical complexity criterion will identify the major problems and evaluate the difficulty of their solution. In addition, this criterion will be used to assess the complexity of software that must be produced for the approach.

4.2.1 Phaseover Approach

The approach of phasing from the DoD architecture to PDN-like architecture requires the solution of technical problems in two major areas. The ability to meet DoD requirements on an X.25 based architecture must be addressed, and the technical plan for phasing the existing system to this new architecture must be developed.

As previously mentioned, the virtual call mechanism of X.25 is limited in the type of service that can be offered, especially in view of the DoD requirements for sophisticated multi-user and transaction oriented applications. It is possible for the DoD to build an enhanced X.25 based network that would support all PDN type services as well as datagram facilities that could be used to provide the other services. The requirements of such a datagram

service could be taken from the existing Internet Protocol by excluding the internetting requirements. These requirements would be met by the X.75 interface.

CCITT has specified a datagram service as part of X.25. This service would make a good starting point for an enhanced DoD specification. Features from IP that would be required in such a specification include reliability, priority, source routing, security marking and guaranteed allocation of resources for a certain level of performance. Features such as source routing are specifically omitted from the X.25 recommendation because they are considered to be in the domain of the PDN supplier. Other features such as reliability and allocation of resources would fit nicely into the options that could be requested from a network.

The other technical problem of using enhanced X.25 networks and compatible protocols is maintaining compatibility with the non-DoD protocols. Once the network moves away from the standard PDN type of configuration, and protocols are added to make use of the enhancements, it is unlikely that they will remain compatible with the comparable protocols on the unenhanced networks. If this incompatibility occurs, then the interoperability is diminished.

Protocols do not have to be symmetrical in order to properly work together, but there will be difficulty designing protocols that will meet the DoD requirements and will also operate with non-DoD counterparts when necessary.

By enhancing the basic network service it would be possible to meet DoD requirements on a PDN type network. The problem of phasing the existing systems and protocols to new ones may be more difficult.

As suggested in the discussion of the phaseover approach, the likely staging plan for integrating new networks and protocols into the IAS would be to remove IP from the architecture and adding an interface between TCP and X.25. This interface would have to provide coordination between the TCP connection management and the X.25 virtual call mechanism. In addition, address mapping would have to be performed between the DoD and CCITT address formats.

Because TCP is designed for use above a datagram support layer (IP), coordination of TCP connections and X.25 virtual calls will not be trivial. TCP's mechanism for opening connections is based on exchanging short messages between the two sites before the connection is actually opened. IP lends itself to this type of usage because of its connectionless nature. With the X.25 support layer, either a virtual call would have to be established before the opening synchronization can occur, or the call would have to be established in conjunction with the opening.

If the call is established before the opening of the TCP connection, then TCP will need some way of communicating to the X.25 interface the need to establish a virtual call to a destination site. After establishing the call or failing to do so, the interface must inform the the TCP module so that the TCP open can take place. Failure of TCP to successfully perform the open, or the closing of the connection after being opened, must also be communicated to the X.25 interface so the virtual call can be cleared.

Establishing the call in conjunction with the TCP open operation requires a similar type of coordination between the TCP protocol and the X.25 interface. In this case the opening TCP synchronization information must be sent as data with the X.25 call request. When the TCP connection does not open, the cause must be determined to be either the inability to establish the X.25 call or the lack of matching SYN information at the remote site.

An alternative strategy is to loosely couple TCP connections with X.25 virtual calls. Instead, the sites would use virtual calls as media by which TCP segments could be sent between hosts. This strategy is similar to the second interoperability approach of using PDNs as network backbones. Taking this approach requires the sites to keep track of which virtual calls have been established, and when they have become inactive.

Although this strategy is simpler, it does not provide any coupling between the transport service (TCP) and the network service (X.25). Establishing the coupling mechanism at this stage will simplify later stages.

After PDN type networks replace the existing DoD backbones by installing the TCP-to-X.25 interface, TCP could be replaced by the transport protocol adopted by non-DoD users. This interface would presumably couple to X.25 more easily than TCP, but would require a new interface to the higher level protocols. The complexity of providing this interface would depend on the similarity between the services provided by the adopted transport protocol and TCP.

The next stage would require the replacement of the higher level protocols with other protocols that are compatible with non-DoD versions. The most visible example of this is the replacement of TELNET or THP with the X.3/X.28/X.29 interface recommendations. Both sets of protocols provide similar service but contain many differences in the details of the service. The difficulty of the problem is a function of the complexity of either mapping old features on the new, incorporating old features in the CCITT mechanism by extension, or removing use of the old features from the applications.

At some point, an enhanced X.25 network may be substituted for the standard PDN type network so that applications specific to the DoD may be supported. In summary, the phaseover approach will present new technical problems at each stage. The number and nature of these problems indicate that the phaseover approach is the most complex.

4.2.2 X.25 Backbone Approach

The problems of using X.25 as the network backbone consist of supporting the interface and performing the required address mapping.

As mentioned in the discussion of this approach, there are three basic techniques that can be used to support the X.25 interface. The most general technique is to allocate the virtual calls based on the need and usage. Any time a packet is received from an IP module at a site, the interface examines the destination address and checks a table to determine if a virtual call is currently set up to that site. If a virtual circuit is not set up, then the

interface attempts to establish one. After it has been established, or if it already exists, the interface uses it for sending the packet. The interface would then update the table concerning the activity of the virtual call. When a virtual call became inactive for a period of time (indicating a cessation of activity between sites) the call would be cleared by the interface.

A more efficient and simpler scheme can be used if the interface is aware that the connected network and all networks on the path to the destination site support datagrams. This scheme provides service which is more consistent with IP's lower level requirements regarding speed and reliability. If the IP only sends to sites on the same physical network then this scheme fits in nicely with the basic approach because it is dependent only on the characteristics of the one network.

When the address space of the interface is restricted to the single physical network, the address mapping with the X.25 backbone approach can be handled by a static table in the interface that translates between the DoD and CCITT address formats. When the interface must consider the use of X.75 interfaces to arbitrary sites on other networks, using a table will not be feasible because of the number of potential destinations. If only a selected number of sites are potential destinations, however, then the table may work for multiple networks.

4.2.3 Protocol Translation Approach

The technical complexity of the protocol translation approach stems from the inherent differences in the basic philosophies of the two architectures. In addition, the technical approach must provide for the actual control of the translator. In this discussion, aspects of the commercial network TELENET will be used as examples of the sort of problems the translator will have to be capable of dealing with.

4.2.3.1 Translator Control

Protocol translation requires two connections to be set up in order to provide interoperability. The first is from the originator's site to the translation site and is established by the originator. The second is from the translation site to the destination site and is established by the translation site on behalf of the originator. Either an explicit or an implicit strategy can be used to establish this second connection.

The explicit strategy requires that the originator establish a connection to the translation site, carry on a conversation with the gateway giving information on the desired destination, and then allowing the gateway to establish the second connection. With this strategy, the originator of the connection is cognizant of the translator's role in providing the connection.

If the implicit strategy is employed, the originator would not be cognizant of the role of the translator. The originator would ask for a connection to the destination site, and this connection request would be sent to the gateway

site. After accepting the connection from the source on behalf of the destination, the gateway would attempt to open a connection to the destination. When the second connection was established, the originator would effectively have a connection from the source to the destination.

Because of its explicit nature, the first strategy requires that extra processing be performed at the originators site to coordinate the establishment of the second connection. This coordination forms a protocol that must be implemented by both the originator and the translator. This protocol would have to be integrated into the protocol architecture at the originator's site.

The implicit strategy requires that the translator accept all connections to the other network. In order to do support this strategy on a PDN, an X.75 type of node must be incorporated into the network. On a DoD type of network, the IP module must be modified to deliver all received packets to the upper level protocol for processing. Normally, the traffic destined for another network would be re-sent by the IP module. After these modifications are made, it will be possible for the translator to extract the information necessary to form the second part of the connection.

4.2.3.2 Addressing

Forming the second connection requires determining the address of the intended destination site in that network's format. The DoD networks use a 32-bit address field with the first 8 bits used as a network identifier, and the last 24 bits used to identify the host. This allows for a network address space of 2^{24} . PDN networks such as TELENET use the X.25 variable length address format. An address is represented as a string of binary coded decimal digits packed two digits per byte. A 12 or 14 decimal digit address is used with the first 4 digits representing a Data Network Identification Code (DNIC), the next 3 digits giving a telephone-like network area code, the next 5 digits specifying the DTE address, and optionally 2 more digits for representing a DTE port. The DNIC contains subfields to identify the continent, country and network to which the rest of the address applies. This allows for a possible address space within a network of 10^{10} .

For a DoD user to identify a TELENET host in a DoD address format, a TELENET network identifier would have to be assigned and used, and the TELENET host number would have to be expressed in the host field. Some mapping of the large address spaces would probably have to be performed, especially because it is not likely that the entire address space will be used.

For a TELENET user to identify a DoD site, a DNIC would have to be assigned for the desired backbone, and the host represented in the DTE address. Again, mapping will have to be done.

Because of the more general nature of the translator, the approach of a table look-up that was applicable in the X.25 network backbone approach may not be applicable to the translator. If the explicit strategy is used, then the addresses can be explicitly provided by the originator.

4.2.3.3 Transport Level Coordination

At the transport level the protocol translator must deal with the issues of synchronizing the two connections and passing flow control information between them.

If an implicit open of the second connection half is to be performed by the translator, then the originator must allow enough time for this to occur. Connections formed by the translator via implicit opens may take longer to complete than connections formed via explicit opens, and originating sites must be aware of this potential delay.

When a connection is opened from an originator on a DoD network, the TCP active open will be received by the translator and will match a previous TCP passive open issued by a higher level protocol (such as TELNET). Without modification to the behavior of the TCP flow control, the open will occur. Data may begin flowing before the translator has a chance to open the connection on the PDN. It is therefore necessary for the open to be made with a window size of zero so that no data will be transmitted on the connection. When the PDN half of the connection is opened, this event must be communicated to the TCP module so that the flow control window can be opened up.

The coordination of flow control is an important issue that must be handled by the translator. Although it is possible for the flow control information to be communicated via the higher level protocols, this approach would be inefficient and require additional buffer space in the translator. The translator must communicate the changing flow control conditions to the other connection side efficiently so that one of the connection does not start to send more data than the other connection is able to accept.

4.2.3.4 Presentation Level Coordination

Because of the developed state of the TELNET protocol and the CCITT virtual terminal support mechanisms, it is worthwhile to examine the translation that must take place between these protocols in some detail to obtain a sense of the technical complexity of the protocol translator approach. In particular the techniques of option negotiation and out-of-band signalling will be examined.

The CCITT virtual terminal support mechanisms consist of the X.3, X.28, and X.29 specifications. X.3 is the specification for the Packet Assembler/Disassembler (PAD) which is the interface between the Start-Stop Mode DTE (i.e., a low speed asynchronous terminal) and packet mode DTEs (i.e., other host computers on the network). The PAD has the basic responsibility for collecting characters received from the terminal into packets and sending them off to the network at appropriate times, and unbundling packets received from the network and delivering them to the terminal. The specific characteristics of these operations are controlled by the values of a set of variables maintained by the PAD for each terminal. These variables, known as PAD parameters, receive default settings by the PAD when a terminal becomes active, and may be changed subsequently by either the terminal or the host

with which the terminal is communicating. The procedure for changing the PAD parameters from the terminal is specified by X.28; from the remote host, X.29 specifies the procedure for changing parameters.

On DoD type networks such as the ARPANET, the TELNET protocol is used for control of terminals. TELNET is generally divided into two halves -- the User and the Server portions. The User TELNET provides the support necessary for controlling the terminals, and is roughly equivalent in functionality to the X.3 and X.28 specifications. The Server TELNET provides the handling at the other end of the connection that interfaces between the network and the process to which the user is communicating. It may be roughly equated to the X.29 specification.

The interface between the terminal handler and the terminal is not affected by the internetting operation, but the interface between the remote host and the terminal handler is. The commands and data exchanged must pass between the two networks, and in doing so, must be translated between the formats and procedures native to each network. This is the basic operation of the translation at the presentation level. The translation done on the data is the repackaging into the format of the destination network type. This is done by the X.29 or TELNET protocol modules as the data is received. The commands received by the modules must be delivered to the virtual terminal translation module for processing. This module would determine how the request can best be translated into the other network's procedures. It would then act as a host process on behalf of the connected host process, or as a terminal on behalf of the connected terminal, and pass the translated request to the other network type's protocol module.

Because the CCITT PAD may take several forms, it is important to clarify which forms of the PAD this translation is directed toward. TELENET supports its own PAD, known as an Interactive Terminal Interface (ITI), which supports a superset of the X.3 defined functions. TELENET may also have individual network subscribers that support exactly X.3 or other variations from the specification. The X.29 procedures that allow remote hosts to set or read PAD parameters are the same for the TELENET PAD as for the X.3 PAD; the procedures allow for the distinguishing of X.3 parameters from network-specific parameters. A list of the TELENET parameters and the X.3 equivalents is given in Figure 7. The intent of the translator will be to handle the full TELENET set of parameters in addition to the X.3 parameters. This will result in the translator being incompatible with individual network subscribers that support their own PADs and utilize PAD parameters that are inconsistent with the TELENET parameters.

Once a virtual terminal connection is established, it can be expected that options will have to be negotiated (from the ARPANET world), or that PAD parameters may have to be set or read (from the TELENET world).

The TELNET protocol provides a method for setting options through the DO, DON'T, WILL, WON'T commands and responses. These options may be initiated from either side of the connection, and at any time during the connection. The TELNET connection starts out with a selection of options that support TELNET-defined "Network Virtual Terminal" Some of the options negotiation for

<u>X.3</u>	<u>TELENET</u>	<u>Description</u>	<u>X.3</u>	<u>TELENET</u>	<u>Description</u>
	0	Natl. Options Marker	30		Abort Output Char.
	1	Linefeed Insertion	31		Intrupt. Proc. Char.
6	2	Net. Message Disp.	32		Automatic Hangup
2	3	Echo	8	33	Flush Output
	4	Echo Mask	34		Transmit on Timers
3	5	Transmit Mask	4	35	Idle Timer
	6	Buffer Size	36		Interval Timer
	7	Command Mask	37		Net. Usage Display
	8	Command Mask	38		Carriage Return Pad
9	9	Carriage Return Pad	39		Padding Options
	10	Linefeed Padding	40		Insert on Break
	11	Tab Padding	5	41	DCE-DTE Flow Control
10	12	Line Width	42		DCE-DTE XON Char.
	13	Page Length	43		DCE-DTE XOFF Char.
	14	Line Folding	44		Generate Break
	15	Reserved	45		APP on Break
	16	Interrupt on Break	46		Input Unlock Option
	17	Break Code	47		2741 Inp. Unl. Timer
	18	NVT Options	48		2741 Inp. Unl. Char.
	19	2741 Int. Key. State	49		2741 Out. Lock Opt.
	20	Half/Full Duplex	50		2741 Out. Lock Timer
	21	Real Character Code	51		2741 Out. Lock Opt.
	22	Printer Style	52		Reserved
	23	Terminal Type	53		Break Options
	24	Permanent Terminal	12	54	DTE-DCE Flow Control
	25	Manual or Automatic	55		DTE-DCE XON Char.
11	26	Rate	56		DTE-DCE XOFF Char.
	27	Delete Character	57		Connection Mode
	28	Cancel Character	1	58	Esc. to Command Mode
	29	Display Character	59		Flush Out. on Break

Figure 7. X.3 and TELENET PAD Parameters

sophisticated options may include sub-negotiation. All sites must support the initial "Network Virtual Terminal" options and may reject any other requests.

X.28 and X.29 procedures allow the setting and reading of PAD parameters and thus control the nature of the service provided by the PAD. The PAD parameters also start out with initial settings that are controlled by the type of PAD being utilized.

When one of these requests is received by a Virtual Terminal level process, it must be passed to the translator to be converted to the closest equivalent on

the other network. If no equivalent exists, then the translator must cause the process to reject the option change. The potential translations from TELNET options to PAD parameters are given below. The PAD parameters not named do not have TELNET equivalents.

0. Binary Transmission -- Maps into PAD parameter 57 (Connection Mode). The TELENET Connection Mode parameter set to "Real" appears to be equivalent to being in binary mode.
1. Echo -- Maps into PAD parameter 3, X.3 parameter 2 (Echo).
2. Reconnection -- No PAD equivalent.
3. Suppress Go Ahead -- Related to PAD parameters 19, 46, 47, 48, 49, 50, and 51. These parameters control the operation of 2741 type terminals for which the Go Ahead signal is intended.
4. Approximate Message Size Negotiation -- Potentially related to PAD parameter 6 (Buffer Size).
5. Status -- No PAD equivalent. *
6. Timing Mark -- No PAD equivalent.
7. Remote Controlled Transmission and Echoing -- No PAD equivalent.
8. Output Line Width -- Maps into PAD parameters 12 (Line Width) and 14 (Line Folding).
9. Output Page Size -- Related to PAD parameter 13 (Page Length).
10. Output Carriage-Return Disposition -- Related to PAD parameters 9, 38 (Carriage Return Padding), 10 (Linefeed Padding), and 39 (Padding Option).
11. Output Horizontal Tabstop -- No PAD equivalent.
12. Output Horizontal Tab Disposition -- Maps into PAD parameter 11 (Tab Padding).
13. Output Formfeed Disposition -- No PAD equivalent.
14. Output Vertical Tabstops -- No PAD equivalent.
15. Output Vertical Tab Disposition -- No PAD equivalent.
16. Output Linefeed Disposition -- Related to PAD parameters 10 (Linefeed Padding) and 39 (Padding Option).
17. Extended ASCII -- No PAD equivalent.

- 18. Logout -- No PAD equivalent.
- 19. Byte Macro -- No PAD equivalent.
- 20. Data Entry Terminal -- No PAD equivalent.
- 21. SUPDUP -- No PAD equivalent.
- 255. Extended-Options-List -- No PAD equivalent.

As can be seen in the previous list there is a high degree of mismatch between the set of options that are available on the two network types. These differences will limit the generality of the interoperability that is available using a translation approach. It should be noted, however, that the features that are most likely to be used for terminal-to-host type communications can be partially translated so that establishing a usable connection should be possible.

Geographically separated users and hosts have made the need for effective out-of-band signaling essential. When this signaling has to be accomplished across multiple networks which employ different mechanisms, the problem becomes more acute. The purpose of out-of-band signaling across a network is to provide the capability of sending a signal over the connection that will bypass other traffic. The the most common objective of this signaling is to interrupt the execution of the user's process. This signaling is also used to allow a process to complete but to discard its output, or to determine that the host is still active. The support of out-of-band signaling is closely tied to the primitives available at the transport level.

TELNET provides its users with a "Synch" mechanism for sending out-of-band signals. The intent of this mechanism is to cause to receiver to scan the input stream for the following out-of-band signals:

- a. Break.
- b. Interrupt Process.
- c. Abort Output.
- d. Are You There.

The TELNET protocol does not require that any specific action take place by the host on receipt of these commands, but suggests (by their naming) what these actions should be.

On a CCITT based network for a terminal-to-host connection, the operation of out-of-band signaling from the terminal is controlled by the setting of the PAD Break Signal Handling parameter (7). If the terminal is connected to the TELENET supported ITI, then additional operations are available from the PAD. These operations are dependent on the setting of TELENET PAD Break Handling parameters (16, 17, 40, 44, 45, 53, and 59) and the Connection Mode parameter (57). Depending on the setting of these parameters, the PAD will perform one

of the following sets of actions on receipt of a break signal from the terminal.

- a. Do nothing.
- b. Escape to Network Command Mode.
- c. Mark the place that the break signal occurred by inserting a Break character in data stream.
- d. Send an X.25 "Reset" packet.
- e. Send an X.25 "Interrupt" packet.
- f. Send an X.25 "Interrupt" packet and discard output to the terminal until signal received from the host.
- g. Send an X.25 "Interrupt" packet and mark the place the break signal occurred by inserting a Break character in the data stream.
- h. Send an X.25 "Interrupt" packet and mark the place the break signal occurred by inserting an Interrupt Process character in the data stream
- i. Send an X.25 "Interrupt" packet, mark the place the break signal occurred by inserting an Abort Output character in the data stream, and discard output to the terminal until signal received from the host.

At the interface between the network and the host, X.29 specifies that an X.25 "Interrupt" packet will be delivered to the host with the data fields set to zero. The specification also provides a procedure by which the host can cause a break signal to be sent to the terminal.

When an out-of-band signal arrives at the protocol translator, it is received by one of the virtual terminal protocol modules and passed to the virtual terminal protocol translator. The translator must then cause a related signal to be issued on the other half of the connection.

Suggested translations from ARPANET to TELENET are as follows:

- a. Break -- cause the X.3/X.28/X.29 module to issue the Break signal as specified by item "g" in the previous list.
- b. Interrupt Process -- cause the X.3/X.28/X.29 module to issue the Interrupt Process signal as specified by item "h" in the previous list.
- c. Abort Output -- cause the X.3/X.28/X.29 module to issue the Abort Process signal as specified by item "i" in the previous list. In addition, discard any output currently queued at the translator for that connection.
- d. Are You There -- There is no TELENET equivalent for this signal.

Translations of out-of-band signals from TELNET to ARPANET are dependent on the the specific signals received by the translator. For the signals listed below the following translations are suggested:

- a. "Interrupt" packet only -- Send a TELNET Break signal.
- b. "Interrupt" packet and Break signal -- Send a TELNET Break signal.
- c. "Interrupt" packet and Interrupt Process signal -- Send a TELNET Interrupt Process signal.
- d. "Interrupt" packet and Abort Output signal -- Send a TELNET Abort Output signal. When new data arrives, the translator must send the signal to the terminal's PAD to stop aborting the output.

Other special function signals that need to be handled by the Translator include the Erase Character and Erase Line signals. These are handled by TELNET as commands, while on TELNET they are part of the character stream. It is therefore necessary for the translator to convert between these two forms. This requires the examination of each character in the TELNET data stream.

Although retention of the exact intended meaning is not always possible in these translations, they should allow a connection to be formed across the two network types and meaningful communication to take place.

4.3 Economic Analysis

The technical features of any proposed approach must be weighed against its anticipated cost over its useful lifetime. Hardware acquisition costs; software design, development, and documentation costs; and ongoing maintenance costs will be covered under this criterion.

4.3.1 Phaseover Approach

The economic costs of doing an architectural phaseover appear high. The proposed plan calls for several stages of software development, with each stage resulting in the elimination of the previous stage. The software to be developed will require a substantial design effort to minimize the impact on existing software. By phasing over to a system based on international standards, commercially available hardware and software can be used for general applications, resulting in potential economic savings over the independent DoD path. Maintenance for commercially available equipment may also be more competitive than for specialized DoD equipment. However, this available hardware and software would likely not meet all government requirements. If it becomes necessary to design enhanced X.25 network hardware, then the costs of this development and the hardware would negate the long term cost advantages.

Another factor is that this approach will affect the entire IAS and all connected sites, resulting in higher costs for installation and documentation than approaches that are more limited in scope.

These apparent high costs must be weighed against the advantage of having a protocol architecture that is compatible with that used by other federal agencies, commercial vendors, and other governments.

4.3.2 X.25 Backbone Approach

The network backbone approach requires the development of the least sophisticated software. The interface called for is necessary only at sites that require PDN connections. These factors should also reduce the cost of software development and installation. The interface should also be relatively simple due to the lack of necessity for changes to other portions of the architectures. The approach may result in savings on hardware acquisition costs if commercially available network hardware can be used and this hardware provides an acceptable level of service. The most serious drawback of this approach is that it may do little to provide the desired interoperability.

4.3.3 Protocol Translation Approach

This approach requires the development of the software necessary to do as complete a protocol translation as possible. The amount of new software required will be a moderate, as discussed in the section on technical complexity. The installation of this software is limited to the translation sites. Depending on whether an implicit or explicit control strategy is used, additional software may be necessary at the site utilizing the protocol translation. Other impacts on existing software may result from the inability of the gateway to translate features that are expected from or required by application programs. Modification of existing application program could entail moderate expense.

4.4 Performance Analysis

The impact of each interconnection approach on the performance that may be obtained from the system will be analyzed using this criterion. Subjects for investigation include potential loss of throughput, delay in opening connections, delay in message delivery, and loss in quality of service.

4.4.1 Phaseover Approach

The impact on performance of taking the phaseover approach will be a function of the type of service required, and the performance that is available from the backbone network used.

For connection oriented services, the impact of the phaseover approach on performance may be minimal. It will be affected by factors that are both network and interface dependent. Network factors include the inherent speed and reliability characteristics, and the number of intra- and inter-network hops that must be taken to support the connection. The interface dependent factors include the overhead necessary to open a connection and transmit data. If the

services offered by X.25 and the transport protocol adopted by the non-DoD community are to be considered equivalent to services offered by TCP and the lower level protocols, comparison of overhead is difficult to make because of the lack of certainty of the details of the non-DoD transport protocol.

Data transfer overhead is a function of the ratio of header length to data length, and the amount of time necessary to process the control information. This processing time will be implementation dependent while the header-to-data ratios are related to traffic size patterns. The likely transport protocol separates the information for sending data from the information for sending acknowledgments. Each of these headers requires 5 octets. Each X.25 packet requires 3 octets of header. Acknowledgments are not necessarily sent with each data unit as in TCP, but if they are, the non-DoD approach would include 16 octets of overhead for each message. By comparison, TCP uses a header whose minimum size is 20 octets. For traffic consisting of short messages (say 4 or fewer characters), the TCP approach might incur slightly higher overhead expense. This expense might be increased if a less active acknowledgment scheme were chosen by the transport protocol. For longer traffic with the same acknowledgment scheme, however, the difference in overhead due to header lengths would be insignificant.

Opening connections in an X.25 based system will require the coordination of the transport protocol with the X.25 virtual call. This coordination is not necessary for the TCP connections using the underlying datagram service and may account for additional overhead in opening and maintaining connections.

The impact of the phaseover approach on the performance of non-connection oriented services is difficult to analyze because of the lack of support in X.25 based systems for this service. However, because of this lack of support, higher overhead can be expected from implementations that provide the same type of service.

4.4.2 X.25 Backbone Approach

IP is intended to utilize a network that provides the capability of efficient delivery of IP packets between network sites. When this network service is replaced by an interface to an X.25 based network, a change in performance can be expected. This change will be the result of two factors: the overhead of the interface between IP and X.25, and the inherent difference of throughput obtainable from the X.25 and datagram based networks.

The interface must select an appropriate virtual call for sending IP packets to their destinations. If a virtual call does not exist to the destination, the interface must establish it. When the call exists as needed, little overhead should be added by the interface, and the packet can be sent immediately assuming the flow control for the call allows it. However, if the call must be set up, there will be delay in the packet delivery while it is established. Other delays will be introduced in the delivery of IP packets between hosts when there is very light traffic between the hosts. This may result in the interface closing the virtual call that exists between the hosts. When the next traffic is generated for that host, the virtual call will have to be

reestablished. Depending on the length of the delay, it is possible for the higher level protocols such as TCP to time out on the original segment, and begin retransmission. This will increase the overhead of the connection support.

If it is possible to utilize the other techniques that were described in the discussion of this approach, such as using datagrams or permanent virtual circuits, then the delay entailed by establishing the X.25 connection will be eliminated. Using datagrams will entail additional overhead for each packet sent because the full X.25 address will be required by the datagram service.

4.4.3 Protocol Translation Approach

The protocol translation approach will have the most severe impact on performance because it requires the most processing of the data and control information between the two network types. This approach requires doing the processing necessary for removing all the protocol layers up to the presentation layer, performing the translation, and then doing the processing for replacing all the protocol layers so that the information can be forwarded to the other network. The approach also requires the examination of each character in the data stream, and the coordination of flow control information between the two connection halves.

In order to translate between the services available on the architectures that are to interoperate, it is necessary for the translation site to act as the termination of each of the connection halves, and do the processing necessary to remove each of the protocol layers up to the presentation level. The data on the connection and the control signals for the presentation device exist at this level. The translating site converts this information into the format of the other architecture as is feasible. It is then necessary to embed the information in the protocols of the other architecture and transmit the data on the other connection. The process of stripping and then embedding protocol layers will increase the time that it takes to transmit information from source to destination.

The status of the flow control information for each end of the connection must be communicated to the other end through the translation site. Because of lag in this communication, the utilization of the internetwork connection may not be as high as when the flow control information goes directly between the ends of the connection.

4.5 Survivability Characteristics

The ability of interconnection schemes to cope with degraded communications services is critical. In particular the impact of single points of failure on critical user connections will be analyzed.

4.5.1 Phaseover Approach

The survivability of the phaseover will be dependent on the topology of the networks employed. Interoperability will allow use of PDNs to provide additional communication paths between different DoD sites. However, connection of the IAS to PDN type networks will permit unauthorized users a more direct form of access to the IAS, which could lead to additional threats that would not exist if the different protocol architectures were employed.

If commercially available equipment were used on the DoD networks, replacement equipment may be more readily available in case of failure. This could reduce the government need to maintain spare equipment.

4.5.2 X.25 Backbone Approach

By using the X.25 backbone approach, the survivability of the IAS in general can be enhanced because it provides the means of alternative connection paths between DoD sites. Links between non-DoD sites supporting DoD protocols and DoD sites are more sensitive to single point of failure concerns because these sites may not have alternative connection paths to the DoD networks. The general reliability of these links will be commensurate with the reliability of commercially available network equipment.

4.5.3 Protocol Translation Approach

The protocol translation approach will have the problem of a single point of failure if there is only one translation site. The failure of the translation site will prevent any interoperability occurring through that site. If multiple sites exist, then this redundancy will increase the survivability of this approach. The survivability is also dependent on the reliability of the PDN equipment and the ability of a site to find a path to connect to the translation site.

4.6 Security Impact

The impact of any interconnection strategy upon network security will be considered. The most significant aspect of this analysis is the application of end-to-end encryption.

4.6.1 Phaseover Approach

By phasing over to a protocol structure that is similar at all sites, it is possible to incorporate end-to-end security in the architecture that will provide all the benefits of this type of protection on a system wide basis. Therefore, if end-to-end security protocols can be properly embedded in the non-DoD architecture, the phaseover approach can make use of this application in order to provide security between the DoD and non-DoD networks.

Another aspect of security that was previously mentioned is that DoD networks will be more vulnerable to malicious users if they are converted to use the same protocols as PDNs, allowing access for interoperability reasons. Such conversion will necessitate a high level of confidence in the security measures employed by the IAS.

4.6.2 X.25 Backbone Approach

The X.25 backbone approach will have little impact on the use of end-to-end security because it does not change the protocol architecture. Thus the current end-to-end technologies will be automatically incorporated if this approach is taken. These technologies can also be used on non-DoD sites that support the DoD protocols to provide end-to-end security on connections to them. End-to-end security measures will prevent text from being in the clear at gateway sites or in X.25 network backbones.

4.6.3 Protocol Translation Approach

Because the translation approach mixes the two architectures, DoD end-to-end security measures cannot be used between two sites on the different network types. If an end-to-end technology is developed for use on X.25 based networks, it could be employed at the translator as part of the processing, but text would then be in the clear at the translation site. The incorporation of security measures in the translator would require that it be trusted to enforce these measures. This requirement would increase the development effort for it.

4.7 Generality for Growth

The interconnection strategy chosen must provide sufficient flexibility to support further growth of the IAS. Use of multicast protocols, interconnection of other networks, addition of new protocols, and incorporation of future technologies will be evaluated.

4.7.1 Phaseover Approach

Because the phaseover approach is based on the the virtual circuit mechanism of the X.25 network, the potential for growth will be limited. The use of an enhanced X.25 network with a datagram facility will make its potential equivalent to that of current DoD technology.

The non-DoD architecture currently lacks a session layer protocol and therefore does not currently incorporate multicast type protocols. Without a datagram service, these types of protocols could be difficult to implement in the X.25 environment.

The addition of more networks to the X.25 architecture should not present any special problems so long as X.75 interfaces are used to connect them. Use of

X.25 on local networks may require a more complex network interface because it would have to perform additional functions. Interfaces to local networks that are broadcast oriented need to be more sophisticated to support the virtual call capability of X.25.

The X.25 virtual call mechanism may not be suitable for use over satellite links which have high bandwidth and long delay characteristics. Special protocols to make use of the satellite characteristics would have to be integrated into an enhanced X.25 network.

4.7.2 X.25 Backbone Approach

Using a datagram protocol as the basis for a protocol architecture is more general than the X.25 virtual call mechanism. Using the X.25 backbone approach will have little effect on this protocol architecture. This approach will therefore not hamper the future growth of the IAS. If performance through the PDN is not as good as that through other DoD backbones, then those protocols that require that performance will suffer.

4.7.3 Protocol Translation Approach

As new protocols are added to the architecture, it will be necessary to incorporate these protocols and their translations into the translator. It is likely that many of these new protocols will not have equivalents on the PDNs and, therefore, no translation will be possible. As mentioned in the discussion of the translation approach, its value will be dependent on its completeness. If more protocol layers are added, and less complete translations are possible, then the upper layers will suffer.

4.8 Transition Impact.

Any changes to the IAS introduce the possibility of user inconvenience while the phaseover is being effected. The impact of such phaseovers will be considered in the evaluation of proposed alternatives.

4.8.1 Phaseover Approach

This approach will have by far the greatest impact during transition. Not only is it the most radical change of all the approaches, but its impact is not limited to a few sites. Instead it affects the entire IAS. Each time a change to the structure is made there will be adjustments to higher level programs to compensate for these changes. These compensations will need to be kept separated from the actual protocol functions so that the compensations and other changes do not become incorporated into the protocols as further development takes place.

4.8.2 X.25 Backbone Approach

By comparison with the previous approach, the X.25 backbone approach will have the least impact on the IAS. It will be incorporated in a limited number of sites and will provide an alternative path for the IP modules in those sites where it is incorporated. It may impact the use of upper level protocols because of the different level of performance that it will offer.

4.8.3 Protocol Translation Approach

As with the previous approaches, the impact of this approach on the IAS is limited to the site where it is implemented. Other sites may be impacted if additional software is necessary for the control of the operation of the gateway site. This approach may also impact the use of upper level protocols because of the different level of performance that it will offer. Upper level protocols should be more noticeably affected by this approach than the X.25 backbone approach because of the likelihood of more serious impact on performance.

5. CONCLUSIONS

The final decision on how interoperability should be approached will be based on the relative weights of the different interoperability objectives. Unless the phaseover to an X.25 type of network is part of the long range planning of the DoD, the first approach should be avoided because of its severe impact on the operations of the IAS, its restrictive structure, and its lack of support for growth without special enhancements.

The second approach, using X.25 networks as backbones, is probably the least complex and is easy to implement, but it will do little to provide interoperability with sites that do not support the DoD protocols.

The third approach of protocol translation will provide a limited form of communication with sites that do not support the DoD protocols, and although it is complex technically, it can be developed and implemented with little impact on the rest of the IAS.

In conclusion, it is recommended that a combination of the last two approaches be taken. The ability to use PDNs as alternative backbones is invaluable for survivability reasons, and it seems feasible for other non-DoD sites on PDNs that have requirements to interoperate with the IAS to support some of the DoD protocols. For the limited set of sites where this support is impossible, the protocol translator could be used to provide the necessary communication capability.

6. REFERENCES

- [1] CCITT Sixth Plenary Assembly Orange Book, Vol. III, Telegraph Technique, Geneva, 1978.
- [2] Defense Communications Agency, Integrated AUTODIN System Architecture Report, December 1977.
- [3] ISO/TC97/SC16, Draft Transport Protocol, N564, December 1980.
- [4] ISO/TC97/SC16, Reference Model of Open Systems Interconnection (Version 4), N227, June 1979.
- [5] Sytek Staff, Preliminary Architecture Report, DCEC Protocols Standardization Program, System Development Corporation TM-7038/200/00, February 1981.